

Listing of Claims

1. (Currently Amended) In a system comprised of a plurality of storage elements, a method for maintaining objects in the storage elements comprising the steps of:

maintaining information regarding which storage elements are storing particular objects in a consistency coordinator which communicates with the storage elements, ~~the consistency coordinator being configured to receive all updates for objects, maintained in the storage elements, from object writers and content providers;~~

responding to a request to update an object by using maintained information to determine which of the storage elements may store stores a copy of the object;

~~targeting the storage elements that include the copy of the object to be updated by:~~

instructing the storage elements, which the consistency coordinator suspects store a copy of the object, to invalidate their copy of the object; and

delaying an updating of the object until it is determined that each storage element instructed to invalidate a copy of the object has either (i) acknowledged that it is not storing a valid copy of the object or (ii) been deemed unresponsive, performing an update of the object after each storage element that includes the copy of the object indicates that the storage element has invalidated the copy of the object or the storage element is determined to be unresponsive—wherein the updating is performed on the object in accordance with one of a plurality of consistency methods where each object has an assigned consistency policy selected on a per-object basis by a content providing application.

2. (Original) The method as recited in claim 1, wherein the step of maintaining information includes maintaining information regarding which storage elements are storing particular objects in the consistency coordinator.

3. (Original) The method as recited in claim 1, wherein the consistency coordinator includes multiple nodes and each node of the consistency coordinator stores information for a different set of objects.

4. (Original) The method as recited in claim 1, wherein the storage elements include at least one cache.

5. (Original) The method as recited in claim 1, wherein the storage elements are included in a distributed system.

6. (Original) The method as recited in claim 1, further comprising the step of obtaining a lock on the object to be updated before performing the update.

7. (Original) The method as recited in claim 1, further comprising the step of sending heart beat messages to obtain availability information about objects from the maintained information to a storage element and from a storage element to the maintained information.

8. (Original) The method as recited in claim 7, further comprising the step of declaring an entity down in response to failing to receive a heart beat.

9. (Original) The method as recited in claim 7, wherein the entity declares itself down in response to failing to receive a heart beat.

10. (Currently Amended) A program storage device readable by machine, tangibly embodying a program of instructions executable by the machine to perform method steps for maintaining strong data consistency, the method steps comprising:

maintaining information regarding which storage elements are storing particular objects in a consistency coordinator which communicates with the storage elements, ~~the consistency coordinator being configured to receive all updates for objects, maintained in the storage elements, from object writers and content providers;~~

responding to a request to update an object by using maintained information to determine which of the storage elements may store stores a copy of the object;

~~targeting the storage elements that include the copy of the object to be updated by:~~

instructing the storage elements, which the consistency coordinator suspects store a copy of the object, to invalidate their copy of the object; and

delaying an updating of the object until it is determined that each storage element instructed to invalidate a copy of the object has either (i) acknowledged that it is not storing a valid copy of the object or (ii) been deemed unresponsive, performing an update of the object after each storage element that includes the copy of the object indicates that the storage element has invalidated the copy of the object or the storage element is

~~determined to be unresponsive wherein the updating is performed on the object in accordance with one of a plurality of consistency methods where each object has an assigned consistency policy selected on a per-object basis by a content-providing application.~~

11. (Currently Amended) In a system comprised of a plurality of storage elements, a method for maintaining stored objects comprising the steps of:

maintaining a consistency coordinator which communicates with the storage elements and stores information regarding which storage elements are storing which objects;

in response to receiving a request to update an object, using information from the consistency coordinator to determine a set of storage elements which may store a copy of the object;

instructing each storage element in the set to invalidate a copy of the object; and
delaying an updating of the object until it is determined that each storage element instructed to invalidate a copy of the object has either (i) acknowledged that it is not storing a valid copy of the object or (ii) been deemed unresponsive.

~~performing the update after each storage element in the set indicates that the storage element has invalidated a copy of the object or the storage element is determined to be unresponsive, wherein the updating is performed on objects in accordance with one of a plurality of consistency methods where each object has an assigned consistency policy selected on a per-object basis by a content-providing application.~~

12. (Original) The method as recited in claim 11, wherein the consistency coordinator includes multiple nodes and further comprising the step of at each node of the consistency coordinator, storing information about which storage elements are storing which objects for a different set of objects.

13. (Original) The method as recited in claim 11, further comprising obtaining a lock from the consistency coordinator by an entity attempting to update an object before performing the update.

14. (Original) The method as recited in claim 11, further comprising the step of sending, from the consistency coordinator to a storage element or from a storage element to the consistency coordinator, heart beat messages to obtain availability information.

15. (Original) The method as recited in claim 14, further comprising an entity expecting a heart beat, declaring itself down in response to failing to receive a heartbeat.

16. (Original) The method as recited in claim 11, wherein the storage elements include at least one cache.

17. (Currently Amended) A program storage device readable by machine, tangibly embodying a program of instructions executable by the machine to perform method steps for maintaining strong data consistency, the method steps comprising:

maintaining a consistency coordinator which communicates with the storage elements

and stores information regarding which storage elements are storing which objects;

in response to receiving a request to update an object, using information from the consistency coordinator to determine a set of storage elements which may store a copy of the object;

instructing each storage element in the set to invalidate a copy of the object; and
delaying an updating of the object until it is determined that each storage element instructed to invalidate a copy of the object has either (i) acknowledged that it is not storing a valid copy of the object or (ii) been deemed unresponsive.

~~performing the update after each storage element in the set indicates that the storage element has invalidated a copy of the object or the storage element is determined to be unresponsive, wherein the updating is performed on objects in accordance with one of a plurality of consistency methods where each object has an assigned consistency policy selected on a per-object basis by a content providing application.~~

18. (Currently Amended) A system for maintaining strong data consistency comprising:
a plurality of storage elements; and
a consistency coordinator configured to receive all updates for objects maintained in the storage elements, from object writers and content providers, the consistency coordinator communicating with the plurality of storage elements and maintaining information about which objects are stored in the plurality of storage elements,

the consistency coordinator providing selective communication to storage elements which include an object to be updated such that for a given object update, the consistency coordinator instructs the storage elements that store a copy of the object to invalidate their copy of the object,

and then delays updating of the object until the consistency coordinator determines that each storage element instructed to invalidate a copy of the object either (i) has acknowledged that it is not storing a valid copy of the object or (ii) is unresponsive. ~~communicates with only these storage elements which include the object to be updated and the consistency coordinator updates objects in accordance with one of a plurality of consistency methods where each object has an assigned consistency policy selected on a per object basis by a content providing application.~~

19. (Original) The system as recited in claim 18, further comprising a writer, which updates the object to be updated.

20. (Original) The system as recited in claim 19, wherein the writer resides on a same node as a storage element.

21. (Original) The system as recited in claim 19, wherein the writer writes an updated object to storage elements after the plurality of storage elements which are to receive the update have invalidated a current copy of the object.

22. (Original) The system as recited in claim 19, wherein the writer writes an updated object to storage elements after the plurality of storage elements which are to receive the update are determined to be unresponsive.

23. (Original) The system as recited in claim 18, further comprising at least one content provider.

24. (Original) The system as recited in claim 23, wherein the content provider resides on a same node as a storage element.

25. (Original) The system as recited in claim 18, further comprising heart beat messages, which may be transmitted between the consistency coordinator and the storage elements to obtain availability information from the consistency coordinator to a storage element or from a storage element to the consistency coordinator.

26. (Original) The system as recited in claim 18, wherein the storage elements include at least one cache.